UC San Diego
JACOBS SCHOOL OF ENGINEERING

Volumetric Segmentation in Electron Microscopy Brain Imaging

Shyam Banuprakash, Ellen Hein, Kim Nguyen, Justin Wahl, Yuping Yu
Advisors: Matthew Madany, Mark Ellisman

Introduction

Advances in volume electron microscopy (VEM) imaging have enabled the accumulation of large-scale, high-resolution biological data. In the neuroscience domain, the potential for these volumes to characterize the dense network of intertwining structures of the brain along with its subcellular components is within reach, but still limited in many ways by the size and complexity of the analysis. We have utilized the high-speed storage and GPU resources provided by the Pacific Research Platform (PRP) and CHASE-CI, managed by the Kubernetes engine, Nautilus, to implement and validate a 3D U-net for multilabel volumetric segmentation trained on a scarcely labeled mouse brain dataset. A deep learning internal zero-shot superresolution was evaluated on downsampled volumes to simulate its effect on pixel classification after x-y resolution boosting. In addition to traditional model performance metrics, a volume rendering tool was created to visualize voxel-level predictions to better understand problematic structures and subcellular features. These models and tools extend the neuroimaging reconstruction framework, NeuroKube.

Problem Statement

Manual labeling of high-resolution VEM brain images is prohibitively time consuming and subject to available expertise in identifying neuronal ultrastructures and organelles. Thus, it becomes necessary to generate models that automate segmentation and do so at scale with high accuracy using very limited training data.

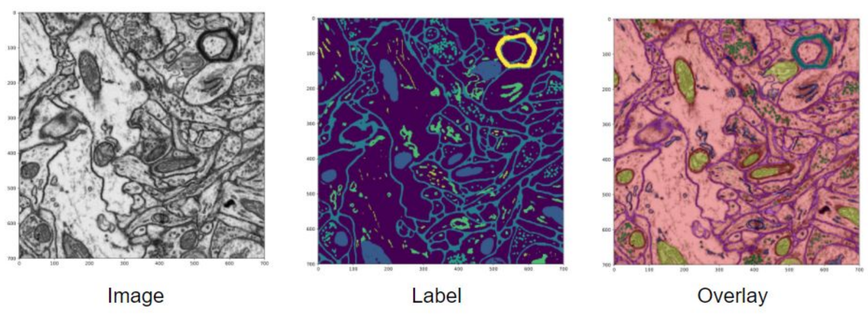

Additionally, not all VEM techniques and preparations are alike and may produce image slices with varying resolutions. Artificial resolution boosting may help resolve smaller subcellular structures, such as neurotransmitter vesicles, that would otherwise be difficult to classify.

Pipeline

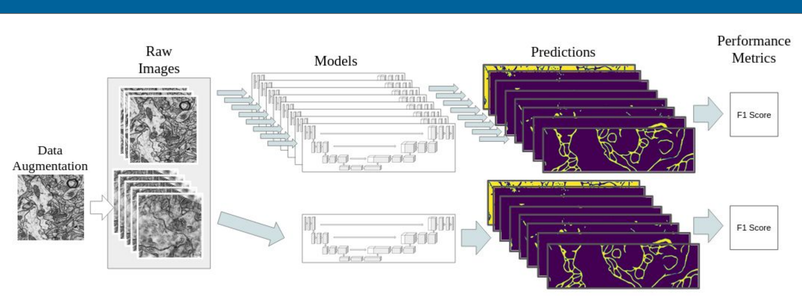

We implemented a 3D U-Net model trained using a small preprocessed 700x700x25 FIB-SEM capture of a mouse cortex. The data was augmented to generate additional samples and add variety to the training set. A multilabel model was trained and evaluated against eight “one-vs-rest” binary models trained individually for each class.

Class
0-cytosol
1-synapses
2-mitochondria
3-membrane
4-neurotransmitter vesicles
5-endoplasmic reticulum
6-microtubules
7-myelin

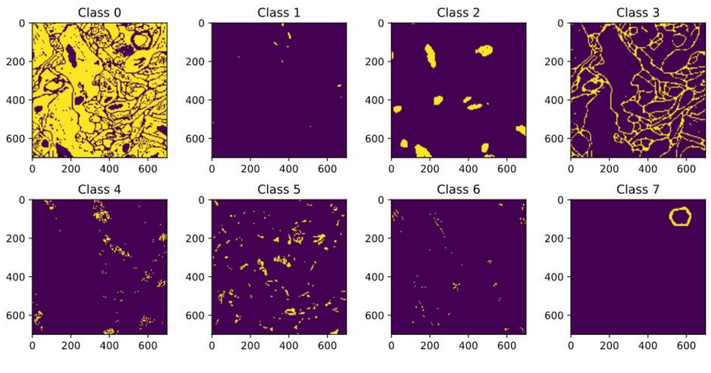

Pipeline

In addition to measuring prediction accuracy, ipyvolume 3D renderings of the confusion matrix was generated to offer a detailed view of misclassified structures.

A zero-shot super resolution (ZSSR) framework using deep internal learning was implemented and evaluated using downsampled 2D slices subjected to artificial resolution boosting. The downsampled and boosted images were then passed through the the trained multilabel U-Net model to observed the effect it had on voxel predictions.

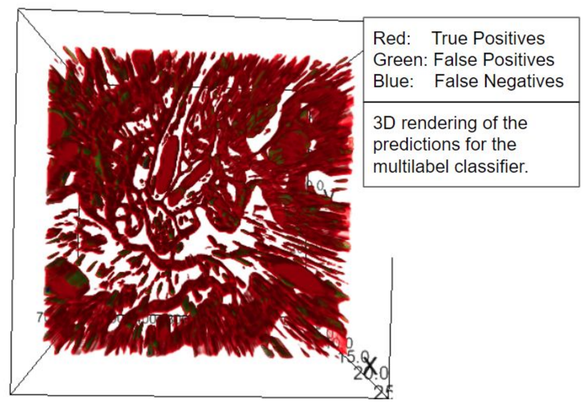

Solution Architecture

Each model was trained in parallel across multiple Jupyter Userlabs within Nautilus, each allocated with a GPU and block storage within the PRP. The models and tools extend images within the ncmir-mm namespace that are accessible through the NeuroKube Gateway. Automatic segmentation can be scaled within Nautilus where the Kubernetes engine will dynamically spin up additional GPUs within the CHASE-CI cyberinfrastructure. A simple webapp built using Streamlit exposes Nautilus resources as a microservice which allows a user to upload zipped VEM stacks that will be passed to the multilabel model and returned as a PyVista volume render.

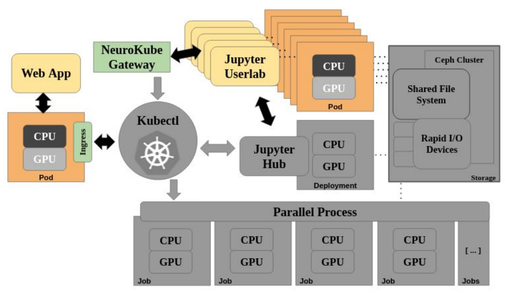

Key Insights

Overall, the multilabel 3D U-Net model performed well with an average weighted F1-score of 92.2 across all evaluated slices. Performance remained the same for many of the classes when compared to its associated binary model predictions, however there did seem to be a slight improvement for classes that are under-represented, such as microtubules and myelin. Applying ZSSR appeared to maintain contrast compared to bilinear scaling, and did improve segmentation slightly across most classes.

Class	% of Total Pixels	F1-Score Binary	F1-Score Multilabel
0-cytosol	73.50%	0.954	0.955
1-synapses	0.29%	0.667	0.666
2-mitochondria	5.78%	0.921	0.922
3-membrane	14.50%	0.879	0.878
4-neurotransmitter vesicles	1.21%	0.752	0.766
5-endoplasmic reticulum	3.25%	0.689	0.711
6-microtubules	0.65%	0.346	0.359
7-myelin	0.73%	0.865	0.904

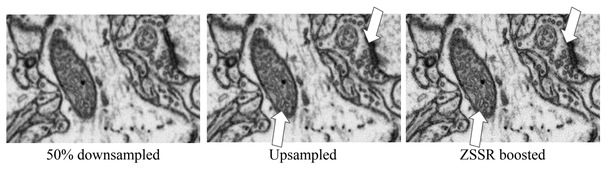

Acknowledgement

We would like to thank our advisors Matthew Madany and Mark Ellisman for providing us with the context and direction for this project. We thank Dr. Ilkay Altintas for providing structure to the capstone process and keeping us on track. We thank the DSE MAS faculty and administrators for providing a seamless transition to remote learning and enabling our continued progress through the COVID-19 pandemic. Last but not least, we would like to thank our families and employers who provided their patience, flexibility, and encouragement during the duration of this two year program.